

# Stream Verification

Justin Thaler\*

## Abstract

We survey models and algorithms for stream verification.

## 1 Problem Definition

Stream verification is concerned with the following setting. A computationally limited client wants to compute some property of a massive input, but lacks the resources to store even a small fraction of the input, and hence cannot perform the desired computation locally. The client therefore accesses a powerful but untrusted service provider (e.g., a commercial cloud computing service), who not only performs the requested computation, but also proves that the answer is correct. An array of closely related models have been introduced to capture this scenario. The following section provides a unified presentation of these models, emphasizing their common features before delineating their differences.

**Stream Verification Model.** Let  $\sigma = \langle a_1, a_2, \dots, a_m \rangle$  be a data stream, where each  $a_i$  comes from a data universe  $\mathcal{U}$  of size  $n$ , and let  $F$  be a function mapping data streams to a finite range  $\mathcal{R}$ . A stream verification protocol for  $F$  involves two parties: a *prover*  $\mathcal{P}$ , and a (randomized) *verifier*  $\mathcal{V}$ . The protocol consists of two stages: a stream observation stage and a proof verification stage.

In the stream observation stage,  $\mathcal{V}$  processes the stream  $\sigma$ , subject to the standard constraints of the data-stream model, i.e., sequential access to  $\sigma$  and limited memory. In the proof verification stage,  $\mathcal{V}$  and  $\mathcal{P}$  exchange a sequence of one or more messages, and afterward  $\mathcal{V}$  outputs a value  $b$ .  $\mathcal{V}$  is allowed to output a special symbol  $\perp$  indicating a rejection of  $\mathcal{P}$ 's claims. Formally,  $\mathcal{V}$  constitutes a stream verification protocol if the following two properties are satisfied.

- **Completeness:** There is some prover strategy  $\mathcal{P}$  such that, for all streams  $\sigma$ , the probability that  $\mathcal{V}$  outputs  $F(\sigma)$  after interacting with  $\mathcal{P}$  is at least  $2/3$ .
- **Soundness:** For all streams  $\sigma$  and all prover strategies  $\mathcal{P}$ , the probability that  $\mathcal{V}$  outputs a value not in  $\{F(x), \perp\}$  after interacting with  $\mathcal{P}$  is at most  $\epsilon \leq 1/3$ .

Here, the probabilities are taken over  $\mathcal{V}$ 's internal randomness. The constants  $2/3$  and  $1/3$  are not essential and are chosen by convention. The parameter  $\epsilon$  is referred to as the *soundness error* of the protocol.

**Costs.** There are five primary costs in any stream verification protocol: (1)  $\mathcal{V}$ 's space usage, (2) the total communication cost, (3)  $\mathcal{V}$ 's runtime, (4)  $\mathcal{P}$ 's runtime, and (5) the number of messages exchanged.

---

\*Yahoo Labs

**Differences Between Models.** There are three primary differences between the various models of stream verification that have been put forth in the literature. The first is whether the soundness condition is required to hold against all cheating provers (such protocols are called *information-theoretically* or *statistically* sound), or only against cheating provers that run in polynomial time (such protocols are called *computationally* sound). The second is the amount and format of the interaction allowed between  $\mathcal{P}$  and  $\mathcal{V}$ . The third is the temporal relationship between the stream observation and proof verification stage — in particular, several models permit  $\mathcal{P}$  and  $\mathcal{V}$  to exchange messages before and during the stream observation stage, and sometimes permit the prover’s messages to depend on parts of the data stream that  $\mathcal{V}$  has not yet seen. In general, more permissive models allow a larger class of problems to be solved efficiently, but may yield protocols that are less realistic.

**Summary of Models.** The *annotated data streaming* (ADS) model, introduced by Chakrabarti et al. [3], is non-interactive:  $\mathcal{P}$  is permitted to send a single message to  $\mathcal{V}$ , with no communication allowed in the reverse direction. Technically, this model permits the contents of  $\mathcal{P}$ ’s message to be interleaved with the stream, in which case each bit of  $\mathcal{P}$ ’s message may be viewed as an “annotation” associated with a particular stream update. However, for most ADS protocols that have developed,  $\mathcal{P}$ ’s message does not need to be interleaved with the stream. Chakrabarti et al. [3] distinguish between two kinds of ADS protocols: *prescient* protocols, in which the annotation sent at any given time can depend on parts of the data stream that  $\mathcal{V}$  has not yet seen, and *online* protocols, which disallow this kind of dependence.

*Streaming interactive proofs* (SIPs) extend the ADS model to allow the prover and verifier to exchange many messages. This model was introduced by Cormode et al. [8].

The *Arthur-Merlin streaming model* was introduced by Gur and Raz [12]: this model is equivalent to a restricted class of SIPs, in which  $\mathcal{V}$  is only allowed to send a single message to  $\mathcal{P}$  (which must consist entirely of random coin tosses, in analogy with the classical complexity class AM), before receiving  $\mathcal{P}$ ’s reply.

The model of *streaming delegation* was introduced by Chung et al. [5], and corresponds to SIPs that only satisfy computational, rather than information-theoretic, soundness.

## 2 Key Results

Obtaining exact answers even for basic problems in the standard data streaming model is impossible using  $o(n)$  space. In contrast, stream verification protocols with  $o(n)$  space and communication costs have been developed for (exactly solving) a wide variety of problems. Many of these protocols have adapted powerful algebraic techniques originally developed in the classical literature on interactive proofs, particularly the *sum-check* protocol of Lund et al. [17]. All of the protocols described in Sections 2.1-2.3 apply even to streams in the strict turnstile update model, where universe items can be deleted as well as inserted. The protocols in Section 2.4 do not support deletions. Unless stated otherwise, all protocols described in this survey are online, i.e., the honest prover’s message at any given time does not depend on parts of the data stream that  $\mathcal{V}$  has not yet seen.

### 2.1 Annotated Data Streams

Chakrabarti et al. [3] showed that prescient ADS protocols can be exponentially more powerful than online ones for some problems. For example, there is a prescient ADS protocol with logarithmic space and communication costs for computing the median of a sequence of numbers:  $\mathcal{P}$  sends  $\mathcal{V}$  the claimed median  $\tau$  at the start of the stream, and while observing the stream,  $\mathcal{V}$  checks that  $|\{j : a_j < \tau\}| \leq m/2$ , and

$|\{j : a_j > \tau\}| \leq m/2$ , which can be done using an  $O(\log m)$ -bit counter. Meanwhile, Chakrabarti et al. [3] proved that any online protocol for MEDIAN with communication cost  $h$  and space cost  $v$  requires  $h \cdot v = \Omega(n)$ , and gave an online ADS protocol achieving these communication–space tradeoffs up to logarithmic factors.

Chakrabarti et al. [3] also gave online ADS protocols achieving identical tradeoffs between space and communication costs for problems including FREQUENCY MOMENTS and FREQUENT ITEMS, and used a lower bound due to Klauck [13] on the *Merlin-Arthur communication complexity* of the SET-DISJOINTNESS function to show that these tradeoffs are optimal for these problems, even among prescient protocols. Subsequent work [7, 23] gave similarly optimal online ADS protocols for several more problems, including maximum matching and counting triangles in graphs, and matrix-vector multiplication. Chakrabarti et al. [2] gave optimized protocols for streams whose length  $m$  is much smaller than the universe size  $n$ .

## 2.2 Streaming Interactive Proofs

Cormode et al. [8] showed that several general protocols from the classical literature on interactive proofs can be simulated in the SIP model. In particular, this includes a powerful, general-purpose protocol due to Goldwasser, Kalai, and Rothblum [10] (henceforth, the GKR protocol). Given any problem computed by an arithmetic or Boolean circuit of polynomial size and polylogarithmic depth, the GKR protocol requires only polylogarithmic space and communication while using polylogarithmic rounds of verifier–prover interaction. This yields SIPs for exactly solving many basic streaming problems with polylogarithmic space and communication costs, including FREQUENCY MOMENTS, FREQUENT ITEMS, and GRAPH CONNECTIVITY. Cormode et al. [8] also gave optimized protocols for specific problems, including FREQUENCY MOMENTS (see the Detailed Example below).

Gur and Raz [12] gave an Arthur-Merlin streaming protocol for the DISTINCT ELEMENTS problem with communication cost  $\tilde{O}(h)$  space cost  $\tilde{O}(v)$  for any  $h, v$  satisfying  $h \cdot v \geq n$ , where the  $\tilde{O}$  notation hides factors that are polylogarithmic in  $n$ . Klauck and Prakash [15] extended this protocol to give an SIP for DISTINCT ELEMENTS with polylogarithmic space and communication costs and logarithmically many rounds of prover–verifier interaction.

Chakrabarti et al. [4] gave *constant-round* online SIPs with logarithmic space and communication costs for many problems, including INDEX, RANGE-COUNTING, and NEAREST-NEIGHBOR SEARCH. These protocols are exponentially more efficient than what can be achieved by constant-round online SIPs in which  $\mathcal{V}$ ’s messages to the prover are independent of the input (such as is required in the Arthur-Merlin streaming model of [12]). For classical interactive proofs where the verifier is not restricted to be streaming, allowing  $\mathcal{V}$ ’s messages to depend on the input does not yield analogous efficiency improvements: Goldwasser and Sipser [11] showed that any interactive proof can be simulated with a polynomial blowup in all costs by an interactive proof in which  $\mathcal{V}$ ’s messages to  $\mathcal{P}$  consist entirely of random coin tosses.

## 2.3 Computationally Sound Protocols

Computationally sound protocols may achieve properties that are unattainable in the information-theoretic setting. For example, they typically achieve reusability, allowing the verifier to use the same randomness to answer many queries. In contrast, most SIPs only support “one-shot” queries, because they require the verifier to reveal secret randomness to the prover over the course of the protocol.

Chung et al. [5] combined the GKR protocol with fully homomorphic encryption (FHE) to give reusable two-message streaming delegation protocols with polylogarithmic space and communication costs for any

problem in the complexity class NC. They also gave reusable four-message protocols with polylogarithmic space and communication costs for any problem in the complexity class P.

Papamanthou et al. [18] gave improved streaming delegation protocols for a class of low-complexity queries including point queries and range search: these protocols avoid the use of FHE, and allow the prover to answer such queries in polylogarithmic time. (In contrast, protocols based on the GKR protocol [5, 8] require the prover to spend time quasilinear in the size of the data stream after receiving a query, even if the answer itself can be computed in sublinear time.) The protocols of [18] are based on *hash trees*.

## 2.4 Models Allowing Linear Total Communication

The literature on stream verification contains two models that depart from those described above in that they allow  $\mathcal{P}$  and  $\mathcal{V}$  to communicate on each stream update. In particular, they allow a linear amount of total communication, but aim to bound the maximum amount of communication exchanged (and processing time spent) during any stream update.

**A Statistically Sound Model.** Klauck and Prakash [14] studied a variant of online SIPs, in which  $\mathcal{P}$  and  $\mathcal{V}$  are allowed to exchange a constant number of messages during each stream update, as long as at most a constant number of machine words are exchanged during each update. They gave protocols in this model for MEDIAN, for determining if a matrix is full-rank, and for approximating the longest-increasing subsequence of the stream.

**A Computationally Sound Model.** Schröder and Schröder [20] introduced a model requiring only computational soundness, which they called the *verifiable data streaming* (VDS) model. This model is targeted at settings in which there are many parties who may wish to verify answers returned by the prover. In the VDS model,  $\mathcal{V}$  is allowed to send a short message to  $\mathcal{P}$  on each stream update, with no communication allowed in the reverse direction. VDS protocols are required to be reusable and publicly verifiable, in the sense that any party who knows  $\mathcal{V}$ 's public key can verify any response by the prover. The VDS model does not permit  $\mathcal{V}$  to update the public key on every stream update, as this would require coordination between all parties wishing to verify information returned by the prover.

There is a trivial VDS protocol for the INDEX problem, requiring polylogarithmic communication on each stream update. Recall that in the INDEX problem, the data stream specifies all entries of a vector  $\mathbf{x} = (x_1, \dots, x_n) \in \{0, 1\}^n$ , followed by an index  $i \in [n]$ , and the goal to output  $x_i$ . In the trivial VDS protocol, for every stream update  $x_j$ ,  $\mathcal{V}$  sends to  $\mathcal{P}$  a digital signature of the tuple  $(j, x_j)$ . After the stream observation phase,  $\mathcal{P}$  sends to  $\mathcal{V}$  the claimed value of  $x_i$  along with a valid digital signature of the tuple  $(i, x_i)$ . The protocol is secure assuming the prover cannot forge valid signatures: under this assumption, the only way  $\mathcal{P}$  can efficiently produce a valid signature of the tuple  $(i, x_i)$  is if  $\mathcal{V}$  had previously sent the signature to  $\mathcal{P}$ .

Schröder and Schröder [20] gave a VDS protocol that supports INDEX queries in a more general setting in which the stream not only contains entries of the vector  $\mathbf{x}$ , but also contains UPDATE operations. Here, an UPDATE operation changes the value of a designated entry  $x_j$  of  $\mathbf{x}$  to a new value  $x'_j$ . However, in the protocol of [20], each UPDATE operation requires bidirectional communication (one message from  $\mathcal{P}$  to  $\mathcal{V}$ , followed by one from  $\mathcal{V}$  to  $\mathcal{P}$ ), as well as an update to  $\mathcal{V}$ 's public key.

Krupp et al. [16] gave related VDS protocols that reduces costs by logarithmic factors.

**Input:**  $\mathcal{V}$  is given oracle access to a  $v$ -variate polynomial  $g$  over finite field  $\mathbb{F}$  and an  $H \in \mathbb{F}$ .

**Goal:** Determine whether  $H = \sum_{(x_1, \dots, x_v) \in \{0,1\}^v} g(x_1, \dots, x_v)$ .

- In the first round,  $\mathcal{P}$  computes the univariate polynomial

$$g_1(X_1) := \sum_{x_2, \dots, x_v \in \{0,1\}^{v-1}} g(X_1, x_2, \dots, x_v),$$

and sends  $g_1$  to  $\mathcal{V}$ .  $\mathcal{V}$  checks that  $g_1$  is a univariate polynomial of degree at most  $\deg_1(g)$ , and that  $H = g_1(0) + g_1(1)$ , rejecting if not.

- $\mathcal{V}$  chooses a random element  $r_1 \in \mathbb{F}$ , and sends  $r_1$  to  $\mathcal{P}$ .
- In the  $j$ th round, for  $1 < j < v$ ,  $\mathcal{P}$  sends to  $\mathcal{V}$  the univariate polynomial

$$g_j(X_j) = \sum_{(x_{j+1}, \dots, x_v) \in \{0,1\}^{v-j}} g(r_1, \dots, r_{j-1}, X_j, x_{j+1}, \dots, x_v).$$

$\mathcal{V}$  checks that  $g_j$  is a univariate polynomial of degree at most  $\deg_j(g)$ , and that  $g_{j-1}(r_{j-1}) = g_j(0) + g_j(1)$ , rejecting if not.

- $\mathcal{V}$  chooses a random element  $r_j \in \mathbb{F}$ , and sends  $r_j$  to  $\mathcal{P}$ .
- In round  $v$ ,  $\mathcal{P}$  sends to  $\mathcal{V}$  the univariate polynomial

$$g_v(X_v) = g(r_1, \dots, r_{v-1}, X_v).$$

$\mathcal{V}$  checks that  $g_v$  is a univariate polynomial of degree at most  $\deg_v(g)$ , rejecting if not.

- $\mathcal{V}$  chooses a random element  $r_v \in \mathbb{F}$  and evaluates  $g(r_1, \dots, r_v)$  with a single oracle query to  $g$ .  $\mathcal{V}$  checks that  $g_v(r_v) = g(r_1, \dots, r_v)$ , rejecting if not.
- If  $\mathcal{V}$  has not yet rejected,  $\mathcal{V}$  halts and accepts.

Figure 1: Description of the sum-check protocol.  $\deg_i(g)$  denotes the degree of  $g$  in the  $i$ th variable.

## 2.5 Implementations

Implementations of the GKR protocol were provided by Cormode et al. in [6] and Thaler in [22]. Cormode et al. [6] also provided optimized implementations of several ADS protocols from [3, 7]. Thaler et al. [24] provided parallelized implementations using Graphics Processing Units.

Qian [19] refined and implemented the streaming delegation protocol for point queries and range search of Papamanthou et al. [18].

Implementations of VDS protocols are described by Schröder and Simkin [21] and Krupp et al. [16].

## 3 Detailed Example

As described in [8], the sum-check protocol can be directly applied to give an SIP for the  $k$ th frequency moment problem with  $\log n$  rounds of prover–verifier iteration, and  $O(\log^2(n))$  space and communication costs. The sum-check protocol is described in Figure 1.

**Properties and Costs of the Sum-check Protocol.** The sum-check protocol satisfies perfect completeness,

and has soundness error  $\epsilon \leq v \cdot \deg(g)/|\mathbb{F}|$ , where  $\deg(g) := \max_{i=1}^v \deg_i(g)$  denotes maximum degree of  $g$  in any variable (see [17] for a proof). There is one round of prover–verifier interaction in the sum-check protocol for each of the  $v$  variables of  $g$ , and the total communication  $O(v \cdot \deg(g))$  field elements.

Note that as described in Figure 1, the sum-check protocol assumes that the verifier has oracle access to  $g$ . However, this will not be the case in applications, as  $g$  will ultimately be a polynomial that depends on the input data stream.

**The SIP for Frequency Moments.** In the  $k$ th frequency moments problem, the goal is to output  $\sum_{i \in \mathcal{U}} f_i^k$ , where  $f_i$  is the number of times item  $i$  appears in the data stream  $\sigma$ . For a vector  $\mathbf{i} = (i_1, \dots, i_{\log n}) \in \{0, 1\}^{\log n}$ , let  $\chi_{\mathbf{i}}(x_1, \dots, x_{\log n}) = \prod_{k=1}^{\log n} \chi_{i_k}(x_k)$ , where  $\chi_0(x_k) = 1 - x_k$  and  $\chi_1(x_k) = x_k$ .  $\chi_{\mathbf{i}}$  is the unique multilinear polynomial that maps  $\mathbf{i} \in \{0, 1\}^{\log n}$  to 1 and all other values in  $\{0, 1\}^{\log n}$  to 0, and it is referred to as the *multilinear extension* of  $\mathbf{i}$ .

For each  $i \in \mathcal{U}$ , associate  $i$  with a vector  $\mathbf{i} \in \{0, 1\}^{\log n}$  in the natural way, and let  $\mathbb{F}$  be a finite field with  $n^k \leq |\mathbb{F}| \leq 4 \cdot n^k$ . Define the polynomial  $\hat{f}: \mathbb{F}^{\log n} \rightarrow \mathbb{F}$  via

$$\hat{f} = \sum_{\mathbf{i} \in \{0, 1\}^{\log n}} f_i \cdot \chi_{\mathbf{i}}. \quad (1)$$

Note that  $\hat{f}$  is the unique multilinear polynomial satisfying the property that  $\hat{f}(\mathbf{i}) = f_i$  for all  $\mathbf{i} \in \{0, 1\}^{\log n}$ .

The  $k$ th frequency moment of  $\sigma$  is equal to  $\sum_{i \in \{0, 1\}^{\log n}} f_i^k = \sum_{\mathbf{i} \in \{0, 1\}^{\log n}} (\hat{f}^k)(\mathbf{i})$ . Hence, in order to compute the  $k$ th frequency moment of  $\sigma$ , it suffices to apply the sum-check protocol to the polynomial  $g = \hat{f}^k$ . This requires  $\log n$  rounds of prover–verifier interaction, and since the total degree of  $\hat{f}^k$  is  $k \cdot \log n$ , the total communication cost is  $O(k \log n)$  field elements, which require  $O(k^2 \log^2 n)$  total bits to specify.

At the end of the sum-check protocol,  $\mathcal{V}$  must compute  $g(r_1, \dots, r_{\log n}) = (\hat{f}^k)(r_1, \dots, r_{\log n})$  for randomly chosen  $(r_1, \dots, r_{\log n}) \in \mathbb{F}^{\log n}$ . It suffices for  $\mathcal{V}$  to evaluate  $z := \hat{f}(r_1, \dots, r_{\log n})$ , since  $(\hat{f}^k)(r_1, \dots, r_{\log n}) = z^k$ . The following lemma establishes that  $\mathcal{V}$  can evaluate  $z$  with a single pass over  $\sigma$ , while storing  $O(\log n)$  field elements.

**Lemma 1.**  $\mathcal{V}$  can compute  $z = \hat{f}(r_1, \dots, r_{\log n})$  with a single streaming pass over  $\sigma$ , while storing  $O(\log n)$  field elements.

*Proof.* Given any stream update  $a_j \in \mathcal{U}$ , let  $\mathbf{a}_j \in \{0, 1\}^{\log n}$  denote the binary vector associated with  $a_j$ . It follows from Equation (1) that  $\hat{f}(r_1, \dots, r_{\log n}) = \sum_{j=1}^m \chi_{\mathbf{a}_j}(r_1, \dots, r_{\log n})$ . Thus,  $\mathcal{V}$  can compute  $\hat{f}(r_1, \dots, r_{\log n})$  incrementally from the raw stream by initializing  $\hat{f}(r_1, \dots, r_{\log n}) \leftarrow 0$ , and processing each update  $a_j$  via:

$$\hat{f}(r_1, \dots, r_{\log n}) \leftarrow \hat{f}(r_1, \dots, r_{\log n}) + \chi_{\mathbf{a}_j}(r_1, \dots, r_{\log n}).$$

$\mathcal{V}$  only needs to store  $(r_1, \dots, r_{\log n})$  and  $\hat{f}(r_1, \dots, r_{\log n})$ , which is  $O(\log n)$  field elements in total.  $\square$

## 4 Open Problems

- Two-message online SIP protocols with logarithmic space and communication costs are known for several functions, including the INDEX function (cf. [4]). It is also known that *existing techniques* cannot yield 2- or 3-message online SIPs of polylogarithmic cost for the MEDIAN or FREQUENCY MOMENTS problems. However, the following is open: exhibit an explicit function  $F: \{0, 1\}^n \rightarrow$



$\{0, 1\}$  that cannot be computed by any online two-message SIP with communication and space costs both bounded above by  $h$ , for some  $h = \omega(\log n)$ . See [4] for details. Candidate functions satisfying this property include SET-DISJOINTNESS and INNER-PRODUCT-MOD-2.

- For several functions  $F: \{0, 1\}^n \rightarrow \{0, 1\}$ , it is known that any online ADS protocol for  $F$  with communication cost  $h$  and space cost  $v$  requires  $h \cdot v = \Omega(n)$ . This lower bound is tight in many cases, such as for the INDEX function (cf. [3]). However, the following is open: exhibit an explicit function that cannot be computed by any online ADS protocol with communication and space costs both bounded above by  $h$ , for some  $h = \omega(n^{1/2})$ .
- Chakrabarti et al. [2] and Thaler [23] have also identified explicit problems which are just as hard for online ADS protocols as they are in the standard data streaming model, in the sense that the sum of  $\mathcal{V}$ 's space cost and the communication cost in any online ADS protocol must be at least as large as the space complexity of a standard streaming algorithm for the problem, up to constant or logarithmic factors. However, it is open to exhibit an explicit function  $F$  that is just as hard in for *prescient* ADS protocols as it is in the standard streaming model.

Formally, the following is open: exhibit an explicit function  $F$  such that, for any constant  $\delta > 0$ ,  $F$  cannot be computed by any prescient ADS protocol with communication and space cost both bounded above by  $O(\text{stream}(F)^{1-\delta})$ . Here,  $\text{stream}(F)$  denotes the minimum space complexity of any streaming algorithm that, for any input stream  $\sigma$ , outputs  $F(\sigma)$  with probability at least  $2/3$ . Candidate functions include GRAPH CONNECTIVITY and GRAPH BIPARTITENESS.

For all three problems, functions satisfying the relevant properties are known to exist by counting arguments. But as stated above, identifying explicit functions satisfying the properties remains open.

## References

- [1] Scott Aaronson and Avi Wigderson. Algebrization: A new barrier in complexity theory. *ACM Trans. Comput. Theory*, 1(1):2:1–2:54, February 2009.
- [2] Amit Chakrabarti, Graham Cormode, Navin Goyal, and Justin Thaler. Annotations for sparse data streams. In Chandra Chekuri, editor, *SODA*, pages 687–706. SIAM, 2014.
- [3] Amit Chakrabarti, Graham Cormode, Andrew McGregor, and Justin Thaler. Annotations in data streams. *ACM Transactions on Algorithms*, 11(1):7, 2014. Preliminary version by the first three authors appeared in *ICALP*, 2009.
- [4] Amit Chakrabarti, Graham Cormode, Andrew McGregor, Justin Thaler, and Suresh Venkatasubramanian. Verifiable stream computation and arthur-merlin communication. In David Zuckerman, editor, *30th Conference on Computational Complexity, CCC 2015, June 17-19, 2015, Portland, Oregon, USA*, volume 33 of *LIPIcs*, pages 217–243. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- [5] Kai-Min Chung, Yael Tauman Kalai, Feng-Hao Liu, and Ran Raz. Memory delegation. In Phillip Rogaway, editor, *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 151–168. Springer, 2011.
- [6] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In Shafi Goldwasser, editor, *ITCS*, pages 90–112. ACM, 2012.

- [7] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Streaming graph computations with a helpful advisor. *Algorithmica*, 65(2):409–442, 2013.
- [8] Graham Cormode, Justin Thaler, and Ke Yi. Verifying computations with streaming interactive proofs. *PVLDB*, 5(1):25–36, 2011.
- [9] Atish Das Sarma, Richard J. Lipton, and Danupon Nanongkai. Best-order streaming model. In *Theory and Applications of Models of Computation*, pages 178–191, 2009.
- [10] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, STOC ’08, pages 113–122, New York, NY, USA, 2008. ACM.
- [11] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In Juris Hartmanis, editor, *STOC*, pages 59–68. ACM, 1986.
- [12] Tom Gur and Ran Raz. Arthur-Merlin streaming complexity. In *Proceedings of the 40th International Colloquium on Automata, Languages and Programming: Part I*, ICALP ’13, Berlin, Heidelberg, 2013. Springer-Verlag.
- [13] Hartmut Klauck. Rectangle size bounds and threshold covers in communication complexity. In *IEEE Conference on Computational Complexity*, pages 118–134, 2003.
- [14] Hartmut Klauck and Ved Prakash. Streaming computations with a loquacious prover. In Robert D. Kleinberg, editor, *ITCS*, pages 305–320. ACM, 2013.
- [15] Hartmut Klauck and Ved Prakash. An improved interactive streaming algorithm for the distinct elements problem. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *ICALP (I)*, volume 8572 of *Lecture Notes in Computer Science*, pages 919–930. Springer, 2014.
- [16] Johannes Krupp, Dominique Schröder, Mark Simkin, Dario Fiore, Giuseppe Ateniese, and Stefan Nürnberger. Nearly optimal verifiable data streaming (full version). *IACR Cryptology ePrint Archive*, 2015:333, 2015.
- [17] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39:859–868, October 1992.
- [18] Charalampos Papamanthou, Elaine Shi, Roberto Tamassia, and Ke Yi. Streaming authenticated data structures. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 353–370. Springer, 2013.
- [19] Yi Qian, Yupeng Zhang, Xi Chen, and Charalampos Papamanthou. Streaming authenticated data structures: Abstraction and implementation. In Gail-Joon Ahn, Alina Oprea, and Reihaneh Safavi-Naini, editors, *Proceedings of the 6th edition of the ACM Workshop on Cloud Computing Security, CCSW ’14, Scottsdale, Arizona, USA, November 7, 2014*, pages 129–139. ACM, 2014.
- [20] Dominique Schröder and Heike Schröder. Verifiable data streaming. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM Conference on Computer and Communications Security*, pages 953–964. ACM, 2012.



- [21] Dominique Schröder and Mark Simkin. Veristream – a framework for verifiable data streaming. In *Financial Cryptography (FC)*, 2015.
- [22] Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In *Proceedings of the 33rd Annual Conference on Advances in Cryptology, CRYPTO’13*, Berlin, Heidelberg, 2013. Springer-Verlag.
- [23] Justin Thaler. Semi-streaming algorithms for annotated graph streams. *Electronic Colloquium on Computational Complexity (ECCC)*, 21:90, 2014.
- [24] Justin Thaler, Mike Roberts, Michael Mitzenmacher, and Hanspeter Pfister. Verifying computations with massively parallel interactive proofs. In *HotCloud*, 2012.
- [25] Victor Vu, Srinath Setty, Andrew J. Blumberg, and Michael Walfish. A hybrid architecture for interactive verifiable computation. In *IEEE Symposium on Security and Privacy*, 2013.